

Remarks

Claim 3 has been re-written in independent form. No amendments affecting claim scope have been made so no new issues arise. It is respectfully submitted that this amendment places the application in better form for appeal because claims 1 and 3 now appear as independent claims and are thus easier to treat separately.

Pursuant to the last amendment filed on February 2, 2008, the applicants were under the impression that the Examiner was prepared to allow the claims. He has now raised a new rejection over Kalampoukas under 35 USC 102(b). As noted in the previous response, in order to meet the test of anticipation, it is essential that

"each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference." *Verdegaal Bros., Inc. v. Union Oil Co.* , 814 F.2d 628, 631, 2 U.S.P.Q.2d 1051, 1053 (Fed. Cir. 1987).

The Federal Circuit has also stated:

"An anticipating reference must describe the patented subject matter with sufficient clarity and detail to establish that the subject matter existed and that its existence was recognized persons of ordinary skill in the field of the invention". *ATD Corp. v. Lydall, Inc.*, 48 USPQ 2d 1321. (Emphasis added)

It is respectfully submitted that the Kalampoukas on a correct reading clearly fails to meet the test of anticipation as established by the courts. The applicants will now explain why this is so. At first sight Kalampoukis appears similar, but after a careful and detailed analysis of the precise teaching of Kalampoukas it will become apparent that

Kalampoukis teaches something quite different from the subject matter claimed. Such an analysis follows.

The Examiner has referred to the passage commencing at col. 6, line 58, which discusses multicast virtual circuits. However, this passage follows on from the discussion of unicast circuits and describes how the Kalampoukas invention may be applied to multicast virtual circuits. Consequently, in order to understand the context of this passage, it is first necessary to fully understand Kalampoukas teaching in relation to unicast circuits.

In the passage commencing at col. 5, line 3, Kalampoukas explains that the rate of transmission of packets belonging to each virtual circuit is controlled at the source based on a feedback mechanism (col. 5, lines 5-7), which is provided by RM cells sent from the source and returned by the destination. There is an RM cell associated with each virtual circuit. The RM cells include a virtual circuit identifier field (see Fig. 2). The source sets the explicit rate (ER) field to the requested value (col. 5, lines 41-42).

As the RM cell traverses the network in the forward direction, the rate allocation algorithm is performed at each port and may modify this field as the RM cell passes through a switch by reducing the value of ER to the maximum amount of bandwidth it is able to allocate on the outgoing link (col. 5, lines 46-47). When the RM cell returns to the source, the latter sets the current rate to the ER value taken from the most recently arrived cell (col 5, line 54) associated with the particular virtual circuit. This is basically the same method as that described in the background of the invention as set forth in the paragraph bridging pages 1 and 2 of the specification.

It is important to note that a separate instance of the rate algorithm, which defines the ER value to be set in the RM cell is, is executed at each output port of the switch (col 5, line

29). It will be noted in Figure 3 that there is a separate rate algorithm processor 30 for each port. Thus, in Figure 1, the algorithm would be executed on the port connected to destination D_2 and on the port connected to switch 4. The passage commencing at col. 6, line 6 further explains how each of the output link modules separates the RM cells and feeds the to the rate allocation algorithm processor 30 for further processing, and how the RM cells, with their updated ER fields, are then combined with the data cells of the virtual connections and transmitted in the forward direction to the destination (not the source). The destination ultimately returns the RM cells, but no processing occurs in the reverse direction in the illustrated embodiment. Upon receipt of the RM cells from the destination, the source is able to set the rate for each virtual circuit as noted at col. 5, line 5.

In the case of multicast virtual circuits, discussed in the paragraph commencing at col. 6, line 58, the same congestion control method is applied. The difference between a multicast connection and a unicast connection is that at a particular switch (switch 3 in Figure 1) a single incoming virtual connection may be split into multiple outgoing virtual connections going to each respective destination. In the multicast example given with reference to Figure 1, the incoming cells from source S_3 at switch 3 are directed to destinations D_2 and switch 4 (for destination D_3). In the forward direction, the system works in the same manner as the unicast case. That is to say a rate calculation algorithm is performed at each outgoing port and the ER field of the outgoing (in the forward direction) RM cells is updated accordingly (see col. 5, line 21, which explains that no processing is performed in the reverse path) before being sent on to the destination.

The passage at col. 7, lines 1 to 5 explains that the current cell rate of a multicast connection is set by the source to the minimum available bandwidth for any of the multicast paths as determined by the collective RM cells received back at the source. The source does not receive RM cells for each multicast flow, but receives collective cells relating to all the multicast flows. The question is, how does Kalampoukis determine this minimum bandwidth? The next passage commencing at col. 7, line 5 is crucial and will be reproduced below. The Examiner's basis for alleged anticipation is based on it, but in the applicant's respectful submission, the Examiner's interpretation of this passage is incorrect.

At any particular switch, the allocation to each one of the RM cells of the multicast connection is the minimum allocation across all flows of the connection. The process of combining the returned RM cells for the several multicast paths can be performed by well known prior art methods. In our preferred embodiment for dealing with multicast flows, we require the switch to find the minimum ER value in the backward direction from among those cells in the multicast flows, that minimum ER value to be then placed in the combined RM cell which is then forwarded in the backward direction (emphasis added).

Thus at switch 3 in the Figure 1, in the forward direction, consistent with the explanation in relation to a unicast connection, a rate calculation is performed at each output port, and the appropriate field of the RM cell for each of the multiple flows updated. One RM cell goes to destination D₂ and one goes to destination D₃ with their respective ER fields updated according to the congestion at the respective output ports on switch 3 (and of course the RM cell going to destination D₃ may also have its ER field updated at switch 4 based on congestion at the output port of switch 4 connected to destination D₃).

In the reverse direction, unlike the unicast case, a common combined cell is formed at each switch where the incoming flow is split into multicast outgoing flows (Switch 3 in Figure 2). In order to determine what value to place in the ER field of the collective cell that is to be returned to the source, the switch 3 then must examine the ER fields of the RM cells returned from destination D₂ and D₃ and pick the one with the lowest value. The important point to note is that the switch 3 requires there to be an ER value in the returning RM cells (traveling in the reverse direction) from each of the flows in order to identify the one with the lowest rate. Thus, an ER calculation must be performed on each flow in the first place in the forward direction in order that the ER value can be examined when the RM cell comes back in the reverse direction in order to determine which ER value to select for placement in the combined cell that is to be returned to the source. In the reverse direction, no processing is performed; the switch 3 merely looks at the incoming ER values in all the flows in the destination and picks the one with the lowest bandwidth for placement in the ER field of the collective RM cell that is returned to the source.

By contrast, claim 1 requires that the slowest stream be first identified, and that once identified the explicit rate calculation be performed only on the slowest stream. Clearly, since the explicit rate calculation is performed only on the slowest stream, it follows as a matter of logic that the slowest stream must be identified by some method other than performing an explicit rate calculation. In the preferred embodiment the identification of the slowest stream is achieved by finding the stream with the greatest phase delay – see page 9, lines 20 -21).

This claimed method is thus quite different from Kalampoukis, who performs the rate calculation on all the streams in the forward direction, and then uses the resulting ER

values in the reverse direction of all the streams to identify the stream with the minimum ER value in the backward direction to be placed in the combined RM cell that is returned to the source. In no way can it be said that Kalampoukis identifies the slowest stream (in the sense discussed above) and then performs the ER calculation only on the slowest stream. This would be impossible in Kalampoukis since he needs to know the ER value in order to determine which stream has the minimum available bandwidth to be included in the combined RM cell to be returned to the source. Without performing the ER calculation he cannot know which stream has the minimum available bandwidth. Clearly, Kalampoukis does not teach performing the ER calculation only on the slowest stream.

Thus, it will be seen that while at first sight the above passage in Kalampoukis seems relevant, on a closer examination Kalampoukis discloses a quite different method, which cannot anticipate claim 1 in accordance with the established tests.

Another important difference is that in Kalampoukis the rate calculation is only performed in the forward direction (col. 5, line 20), so the result of the ER calculation is sent to the destination, not the source as required by claim 1. While the destination ultimately returns this ER value in the RM cell back toward the switch in the reverse direction (and this ER value may or may not reach the source depending on whether it is picked during the process to form the combined RM cell), sending the RM cell to the destination could not reasonably be regarded as satisfying the limitation “transmitting a result of the slowest stream ER calculation back to the source”, especially bearing in mind that the claim has to be interpreted in the light of the disclosure, which explains in the passage bridging pages 7 and 8 how the locally computed ER value is compared with the ER value in the counter-flowing RM cell, and if the former is less than the latter, the locally generated ER value is used to update the value in the RM cell.

By analogy if the Examiner were driving his car and upon arriving at an intersection was directed by a police officer to turn left toward a specific destination, but instead decided to turn right, the police officer would not accept an explanation that the Examiner intended to turn right, drive to the end of the road, and do a U-turn, before heading back in the left direction toward the ultimate destination. In other words, a limitation that the result of the ER calculation is transmitted “back to the source” cannot be satisfied by a teaching that the computed ER value is transmitted to the destination regardless of what happens subsequently.

Moreover, the subject matter of claim 1 is not obvious under 35 USC 103 because the object of the invention is precisely to eliminate the need to perform an ER calculation on each port as taught by Kalampoukis. It is completely contrary to the teaching of Kalampoukis to avoid performing such multiple calculations, and it is by no means apparent how this could be achieved within the context of the Kalampoukis teaching. Unlike the invention Kalampoukis has to perform an ER calculation on each flow.

With regard to claim 3, as noted above, it follows from the wording of claim 1 that, unlike Kalampoukis, in the present invention the step of identifying the slowest stream is not achieved by using the results of the ER calculation (It will be recalled that Kalampoukis identifies the stream with the minimum bandwidth simply by looking at the ER values in the RM cells being returned to the source from the destination). Claim 3 recites a preferred method of identifying the slowest stream (which as noted takes place before the ER calculation is performed), wherein the incoming cells are stored in a linked list, a read pointer is maintained for each stream, and the slowest stream is identified by identifying the read pointer associated with the earliest cell in the linked list. The


Examiner's assertion that Kalampoukis teaches this is completely baffling because Kalampoukis uses the calculated ER values to find the stream with minimum bandwidth.

In support of the Examiner's allegation that Kalampoukis teaches such a method, the Examiner refers to Figures 3 and 4, which show a buffered switch, and col. 6 lines 34 - 45). This passage says nothing about identifying the slowest stream from the read pointer. On the contrary, each output port is associated with a rate processor, which implements the rate algorithm to obtain the ER value. This is completely contrary to the method of identifying the slowest stream as set forth in claim 3. Nowhere does Kalampoukis teach that the slowest stream is identified from the read pointer associated with the earliest cell in a linked list.

Where is the temporarily ordered linked list? The Examiner refers to col. 5, lines 25 -31. Firstly, this passage refers to unicast virtual connections and says nothing about a linked list relating the buffers, but more importantly this passage states that separate instances of the rate algorithm, which is the algorithm for determining the ER value, are executed at each port. In the applicant's respectful submission, the Examiner has taken the wording from the applicant's claim and attributed it to certain passages in Kalampoukis that simply do not support such an interpretation.

Finally, with regard to the limitation that read pointer is maintained for each stream to provide an index into the linked list and that the slowest stream is identified by identifying the read pointer associated with the temporally earliest cell in the linked list, the Examiner merely refers to the fact that the RM cell has a VC identifier and somehow seems to read into this the entire claim limitation. Col. 5, lines 50 -63 merely explains that the source uses the ER field of the returned RM cell to set the current outgoing rate in

unicast connections. Col. 7, lines 30 – 44 merely describe the structure of the RM cell. The read pointer relates to the location in memory of the next cell to be read out of memory. It has nothing to do with the VCI field. In the applicant's respectful submission, this aspect of the Examiner's rejection simply makes no sense.



What the invention is saying is that in the above diagram, the incoming stream is split into flows 2, 3, 4 associated with buffers 5, 6, 7 of the respective multicast streams, and that the slowest stream is identified by looking at the read pointers of the outgoing buffers (because the slowest stream will be the one with the temporally earliest cell), and that this stream is then used to perform the rate calculation to find the ER value which is sent back to the source.

By contrast in Kalampoukis, an ER calculation is performed on each of the outgoing buffers without regard to which is the slowest, this ER value is placed in the RM cells that are sent to the destination, not the source, and when the RM cells come back from the source, the switch looks at the ER values in the RM cells from each stream, notes which is the minimum ER value, and places this in the ER field of the combined cell that is eventually returned to the source. Such an arrangement is demonstrably different from the invention as claimed.

In the applicant's submission, the Examiner has taken isolated passages of Kalampoukis out of context and read into them far more than is actually stated. Indeed, in the applicant's respectful submission, the Examiner has made the fundamental error of interpreting vague passages in Kalampoukis in the light of the applicant's teachings rather than in the light of Kalampoukis' own teachings. On a detailed review it is clear that the mode of operation of Kalampoukis is fundamentally different from the invention as claimed for at least the reasons noted above, and indeed is essentially the same as the admitted prior art that the invention seeks to improve upon. At the very least Kalampoukis clearly fails to meet the rigorous test of anticipation under 35 USC 102(b). For the reasons given above, the applicants believe that the independent claims in their present form are clearly distinguished over Kalampoukis, which discloses a system that is quite different in operation. If the Examiner would prefer any amendments to the wording which he might resolve any perceived ambiguities, he is kindly requested to telephone the undersigned to avoid the need for an appeal.

It is believed that the application is in condition for allowance for essentially the same reasons as before. Allowance and reconsideration are therefore earnestly solicited.

Respectfully submitted,



Registration No. 34519
Richard J. Mitchell
Agent of Record

MARKS & CLERK
P. O. Box 957, Station B,
Ottawa, Ontario, Canada
K1P 5S7
(613) 236-9561